arXiv:1001.1470v1 [cs.DS] 10 Jan 2010

# A New Approximation Technique for Resource-Allocation Problems

Barna Saha[1] * Aravind Srinivasan[1] †
[1]University of Maryland, College Park, MD 20742, USA
barna@cs.umd.edu srin@cs.umd.edu

**Abstract:** We develop a rounding method based on random walks in polytopes, which leads to improved approximation algorithms and integrality gaps for several assignment problems that arise in resource allocation and scheduling. In particular, it generalizes the work of Shmoys & Tardos on the generalized assignment problem in two different directions, where the machines have hard capacities, and where some jobs can be dropped. We also outline possible applications and connections of this methodology to discrepancy theory and iterated rounding.

**Keywords:** Scheduling; rounding; approximation algorithms; integrality gap

## 1 Introduction

The "relax-and-round" paradigm is a well-known approach in combinatorial optimization. Given an instance of an optimization problem, we *enlarge* the set of feasible solutions $I$ to some set $I' \supset I$ – often the linear-programming (LP) relaxation of the problem; we then map an (efficiently computed, optimal) solution $x^* \in I'$ to some "nearby" $x \in I$ and prove that $x$ is near-optimal in $I$. This second "rounding" step is often a crucial ingredient, and many general techniques have been developed for it. In this work, we present a new rounding methodology which leads to several improved approximation algorithms in scheduling, and which, as we explain, appears to have connections and applications to other techniques and problems, respectively.

We next present background on (randomized) rounding and a fundamental scheduling problem, before describing our contribution.

Our work generalizes various *dependent* randomized rounding techniques that have been developed over the past decade or so. Recall that in randomized rounding, we use randomization to map $x^* = (x_1^*, x_2^*, \ldots, x_n^*)$ back to some $x = (x_1, x_2, \ldots, x_n)$ [39]. Typically, we choose a value $\alpha$ that is problem-specific, and, *independently* for each $i$, define $x_i$ to be 1 with probability $\alpha x_i^*$, and to be 0 with the complementary probability of $1 - \alpha x_i^*$. Independence can, however, lead to noticeable deviations from the mean for random variables that are *required* to be very close to (or even be equal to) their mean. A fruitful idea developed in [27, 32, 45] is to carefully introduce *dependencies* into the rounding process: in particular, some sums of random variables are held fixed with probability one, while still retaining randomness in the individual variables and guaranteeing certain types of negative-correlation properties among them. See [1] for a related deterministic approach that precedes these works. These dependent-rounding approaches lead to numerous improved approximation algorithms in scheduling and packet-routing [1, 27, 32, 45].

We now introduce a fundamental scheduling model, which has spurred many advances and applications in combinatorial optimization, including linear-, quadratic- & convex-programming relaxations and new rounding approaches [6, 8, 10, 15, 21, 29, 32, 34, 41, 43]. This model, *scheduling with unrelated parallel machines* (UPM) – and its relatives – play a key role in this work. Herein, we are given a set $J$ of $n$ jobs, a set $M$ of $m$ machines,

*Dept. of Computer Science, University of Maryland, College Park, MD 20742. Research supported by NSF Award NSF CCF-0728839.

†Dept. of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. Supported in part by NSF ITR Award CNS-0426683 and NSF Award CNS-0626636.

and non-negative values $p_{i,j}$ ($i \in M,\ j \in J$): each job $j$ has to be assigned to some machine, and assigning it to machine $i$ will impose a processing time of $p_{i,j}$ on machine $i$. (The word "unrelated" arises from the fact that there may be no pattern among the given numbers $p_{i,j}$.) Variants such as the type of objective function(s) to be optimized in such an assignment, whether there is an additional "cost-function", whether a few jobs can be dropped, and situations where there are release dates for, and precedence constraints among, the jobs, lead to a rich spectrum of problems and techniques. We now briefly discuss two such highly-impactful results [34, 41]. The primary UPM objective in these works is to minimize the *makespan* – the maximum total load on any machine. It is shown in [34] that this problem can be approximated to within a factor of 2; furthermore, even some natural special cases cannot be approximated better than $1.5$ unless $P = NP$ [34]. Despite much effort, these bounds have not been improved. The work of [41] builds on the upper-bound of [34] to consider the *generalized assignment problem* (GAP) where we incur a cost $c_{i,j}$ if we schedule job $j$ on machine $i$; a simultaneous $(2, 1)$–approximation for the (makespan, total cost)-pair is developed in [41], leading to numerous applications (see, e.g., [2, 17]).

We generalize the methods of [1, 27, 31, 32, 45], via a type of random walk toward a vertex of the underlying polytope that we outline next. We then present several applications in scheduling and bipartite matching through problem-specific specializations of this approach, and discuss further prospects for this methodology.

The rounding approaches of [1, 27, 31, 45] are generalized to linear systems as follows in [32]. Suppose we have an $n$-dimensional constraint system $Ax \leq b$ with the additional constraints that $x \in [0, 1]^n$. This will often be a LP-relaxation, which we aim to round to some $y \in \{0, 1\}^n$ such that some constraints in "$Ay \leq b$" hold with probability one, while the rest are violated "a little" (with high probability). Given some $x \in [0, 1]^n$, the rounding approach of [32] is as follows. First, we assume without loss of generality that $x \in (0, 1)^n$: those $x_j$ that get rounded to 0 or 1 at some point, are held fixed from then on. Next, we "judiciously" drop some of the constraints in "$Ax \leq b$" until the number of constraints becomes smaller than $n$, thus making the system linearly-dependent – leading to the efficient computation of an $r \in \Re^n$ that is in the nullspace of this reduced system. We then compute positive scalars $\alpha$ and $\beta$ such that $x_1 := x + \alpha r$ and $x_2 := x - \beta r$ both lie in $[0, 1]^n$, and both have at least one component lying in $\{0, 1\}$; we then update $x$ to a random $Y$ as: $Y := x_1$ with probability $\beta/(\alpha + \beta)$, and $Y := x_2$ with the complementary probability $\alpha/(\alpha + \beta)$. Thus we have rounded at least one further component of $x$, and also have the useful property that for all $j$, $\mathbf{E}[Y_j] = x_j$. Different ways of conducting the "judicious" reduction lead to a variety of improved scheduling algorithms in [32]. The setting of [27, 45] on bipartite $b$-matchings can be interpreted in this framework.

We further generalize the above-sketched approach of [32]. Suppose we are given a polytope $\mathcal{P}$ in $n$ dimensions, and a *non-vertex* point $x$ belonging to $\mathcal{P}$. An appropriate basic-feasible solution will of course lead us to a vertex of $\mathcal{P}$, but we approach (not necessarily reach) a vertex of $\mathcal{P}$ by a random walk as follows. Let $\mathcal{C}$ denote the set of constraints defining $\mathcal{P}$ which are satisfied *tightly* (i.e., with equality) by $x$. Then, note that there is a non-empty linear subspace $S$ of $\Re^n$ such that for any nonzero $r \in S$, we can travel up to some strictly-positive distance $f(r)$ along $r$ starting from $x$, while staying in $\mathcal{P}$ and *continuing* to satisfy all constraints in $\mathcal{C}$ tightly. Our broad approach to conduct a random move $Y := x + R$ by choosing an appropriately random $R$ from $S$, such that the property "$\mathbf{E}[Y_j] = x_j$" of the previous paragraph still holds. In particular, let **RandMove**$(x, \mathcal{P})$ – or simply **RandMove**$(x)$ if $\mathcal{P}$ is understood – be as follows. Choose a nonzero $r \in S$ arbitrarily, and set $Y := x + f(r)r$ with probability $f(-r)/(f(r)+f(-r))$, and $Y := x - f(-r)r$ with the complementary probability of $f(r)/(f(r) + f(-r))$. Note that if we repeat **RandMove**, we obtain a random walk that finally leads us to a vertex of $\mathcal{P}$; the high-level idea is to intersperse this walk with the idea of "judiciously dropping some constraints" from the previous paragraph, as well as combining certain constraints together into one. Three major differences from [32] are: (a) the care given to the tight constraints $\mathcal{C}$, (b) the choice of which constraint to

drop being based on $\mathcal{C}$, and (c) clubbing some constraints into one. As discussed next, this recipe appears fruitful in a number of directions in scheduling, and as a new rounding technique in general.

*Capacity constraints on machines, random matchings with sharp tail bounds.* Handling "hard capacities" – those that cannot be violated – is generally tricky in various settings, including facility-location and other covering problems [19, 26, 36]. Motivated by problems in crew-scheduling [22, 40] and by the fact that servers have a limit on how many jobs can be assigned to them, the natural question of scheduling with a hard capacity-constraint of "at most $b_i$ jobs to be scheduled on each machine $i$" has been studied in [18, 48, 50–52]. Most recently, the work of [18] has shown that this problem can be approximated to within a factor of 3 in the special case where the machines are *identical* (job $j$ has processing time $p_j$ on any machine). In § 2, we use our random-walk approach to generalize this to the setting of GAP and obtain the GAP bounds of [41] – i.e., approximation ratios of 2 and 1 for the makespan and cost respectively, while satisfying the capacity constraints: the improvements are in the more-general scheduling model, adding the cost constraint, and in the approximation ratio. We anticipate that such a capacity-sensitive generalization of [41] would lead to improved approximation algorithms for several applications of GAP, and present one such in Section 5.

Theorem 1 generalizes such capacitated problems to random bipartite ($b$-)matchings with target degree bounds and sharp tail bounds for given linear functions; see [23] to applications to models for complex networks. Recall that a ($b$)-matching is a subgraph in which every vertex $v$ has degree at most $b(v)$. Given a *fractional* ($b$)-matching $x$ in a bipartite graph $G = (J, M, E)$ of $N$ vertices and a collection of $k$ linear functions $\{f_i\}$ of $x$, many works have considered the problem of constructing ($b$-)matchings $X$ such that $f_i(X)$ is "close" to $f_i(x)$ simultaneously for each $i$ [3, 27, 28, 38]. The works [28, 38] focus on the case of constant $k$; those of [3, 27] consider general $k$, and require the usual "discrepancy" term of $\Omega(\sqrt{f_i(x) \log N})$ in $|f_i(X) - f_i(x)|$ for most/all $i$; in a few cases, $o(N)$ vertices will have to remain unmatched also. In contrast, Theorem 1 shows that if there is one structured objective function $f_i$ with bounded coefficients associated with each $i \in M$, then in fact all the $|f_i(X) - f_i(x)|$ can be bounded independent of $N$. This appears to be the first such result here, and helps with equitable max-min fair allocations as discussed below.

*Scheduling with outliers: makespan and fairness.* Note that the $(2, 1)$ bicriteria approximation that we obtain for GAP above, generalizes the results of [41]. We now present such a generalization in another direction: that of "outliers" in scheduling [29]. For instance, suppose in the "processing times $p_{i,j}$ and costs $c_{i,j}$" setting of GAP, we also have a profit $\pi_j$ for choosing to schedule each job $j$. Given a "hard" target profit $\Pi$, target makespan $T$ and total cost $C$, the LP-rounding method of [29] either proves that these targets are not simultaneously achievable, or constructs a schedule with values $(\Pi, 3T, C(1 + \epsilon))$ for any constant $\epsilon > 0$. Using our rounding approach, we improve this to $(\Pi, (2 + \epsilon)T, C(1 + \epsilon))$ in § 3. (The factors of $\epsilon$ in the cost are required due to the hardness of knapsack [29].) Also, fairness is a fundamental issue in dealing with outliers: e.g., in repeated runs of such algorithms, we may not desire long starvation of individual job(s) in sacrifice to a global objective function. Theorem 7 accommodates fairness in the form of scheduling-probabilities for the jobs that can be part of the input.

*Max-Min Fair Allocation.* This problem, also known as the *Santa Claus* problem, is the max-min version of UPM, where we aim to maximize the minimum "load" (viewed as utility) on the machines; it has received a good deal of attention recently [4, 5, 8, 10, 15, 24]. We are able to employ dependent randomized rounding to near-optimally determine the integrality gap of a well-studied LP relaxation. Also, Theorem 1 lets us generalize a result of [14] on max-min fairness to the setting of equitable partitioning of the jobs; see § 4.

**Directions for the future: some potential connections and applications.** Distributions on *structured* matchings in bipartite graphs is a topic that models many scenarios in discrete optimization, and we view our work as a useful contribution to it. We explore further applications and connections in § 6. A general question involving "rounding well" is the *lattice approximation* problem [39]: given $A \in \{0, 1\}^{m \times n}$ and $p \in [0, 1]^n$,

we want a $q \in \{0,1\}^n$ such that $\|A \cdot (q-p)\|_\infty$ is "small"; the linear discrepancy of $A$ is defined to be $\text{lindisc}(A) = \max_{p \in [0,1]^n} \min_{q \in \{0,1\}^n} \|A \cdot (q-p)\|_\infty$. The field of *combinatorial discrepancy theory* [13] has developed several classical results that bound $\text{lindisc}(A)$ for various matrix families $A$; column-sparse matrices have received much attention in this regard. Section 6 discusses a concrete approach to use our method for the famous Beck-Fiala conjecture on the discrepancy of column-sparse matrices [12], in the setting of random matrices. § 6 also suggests that there may be deeper connections to *iterated rounding*, a fruitful approach in approximation algorithms [25, 30, 33, 42, 49]. We view our approach as having broader connections/applications (e.g., to open problems including capacitated facility location [36]), and are studying these directions.

## 2 Random Matchings with Linear Constraints, and GAP with Capacity Constraints

We develop an efficient scheme to generate random subgraphs of bipartite graphs that satisfy hard degree-constraints and near-optimally satisfy a collection of linear constraints:

THEOREM 1 *Let $G = (J, M, E)$ be a bipartite graph with "jobs" $J$ and "machines" $M$. Let $\mathcal{F}$ be the collection of edge-indexed vectors $y$ (with $y_{i,j}$ denoting $y_e$ where $e = (i,j) \in E$). Suppose we are given: (i) an integer* requirement *$r_j$ for each $j \in J$ and an integer* capacity *$b_i$ for each $i \in M$; (ii) for each $i \in M$, a linear objective function $f_i : \mathcal{F} \to \Re$ given by $f_i(y) = \sum_{j:\ (i,j) \in E} p_{i,j} y_{i,j}$ such that $0 \le p_{i,j} \le \ell_i$ for each $j$, and (iii) a vector $x \in \mathcal{F}$ with $x_e \in [0,1]$ for each $e$. Then, we can efficiently construct a random subgraph of $G$ given by a binary vector $X \in \mathcal{F}$, such that: (a) with probability one, each $j \in J$ has degree at least $r_j$, each $i \in M$ has degree at most $b_i$, and $|f_i(X) - f_i(x)| < \ell_i \ \forall i$; and (b) for all $e \in E$, $\mathsf{E}[X_e] = x_e$.*

We will now prove an important special case of Theorem 1: GAP with individual capacity constraints on each machine. This special case captures much of the essence of Theorem 1; the full proof of Theorem 1 is deferred to the final version of this work. The capacity constraint specifies the maximum number of jobs that can be scheduled on any machine, and is a hard constraint. Formally the problem is as follows, where $x_{i,j}$ is the indicator variable for job $j$ being scheduled on machine $i$. Given $m$ machines and $n$ jobs, where job $j$ requires a processing time of $p_{i,j}$ in machine $i$ and incurs a cost of $c_{i,j}$ if assigned to $i$, the goal is to minimize the makespan $T = \max_i \sum_j x_{i,j} p_{i,j}$, subject to the constraint that the total cost $\sum_{i,j} x_{i,j} c_{i,j}$ is at most $C$ and for each machine $i$, $\sum_j x_{i,j} \le b_i$. $C$ is the given upper bound on total cost and $b_i$ is the capacity of machine $i$, that must be obeyed.

Our main contribution here is an efficient algorithm **Sched-Cap** that has the following guarantee, generalizing the GAP bounds of [41]:

THEOREM 2 *There is an efficient algorithm **Sched-Cap** that returns a schedule maintaining all the capacity constraints, of cost at most $C$ and makespan at most $2T$, where $T$ is the optimal makespan with cost $C$ that satisfies the capacity constraints.*

We guess the optimum makespan $T$ by binary search as in [34]. If $p_{i,j} > T$, $x_{i,j}$ is set to 0. The solution to the following integer program gives the optimum schedule:

$$\sum_{i,j} c_{i,j} x_{i,j} \le C \qquad \text{(Cost)}$$

$$\sum_{i,j} x_{i,j} = 1 \ \forall j \qquad \text{(Assign)}$$

$$\sum_{j} p_{i,j} x_{i,j} \le T \ \forall i \qquad \text{(Load)}$$

$$\sum_{j} x_{i,j} \le b_i \ \forall i \qquad \text{(Capacity)}$$

$$x_{i,j} \in \{0,1\} \ \forall i,j$$

$$x_{i,j} = 0 \quad \text{if } p_{i,j} > T$$

We relax the constraint "$x_{i,j} \in \{0,1\}\ \forall (i,j)$" to "$x_{i,j} \in [0,1]\ \forall (i,j)$" to obtain the LP relaxation **LP-Cap**. We solve the LP to obtain the optimum LP solution $x^*$; we next show how **Sched-Cap** rounds $x^*$ to obtain an integral solution within the approximation guarantee.

Note that $x^*_{i,j} \in [0,1]$ denotes the "fraction" of job $j$ assigned to machine $i$. Initialize $X = x^*$. The algorithm is composed of several iterations.

The random value of the assignment-vector $X$ at the end of iteration $h$ of the overall algorithm is denoted by $X^h$. Each iteration $h$ conducts a randomized update using the **RandMove** on the polytope of a linear system constructed from a *subset* of the constraints of **LP-Cap**. Therefore, by induction on $h$, we will have for all $(i, j, h)$ that $\mathsf{E}\left[X^h_{i,j}\right] = x^*_{i,j}$.

Let $J$ and $M$ denote the set of jobs and machines, respectively. Suppose we are at the beginning of some iteration $(h+1)$ of the overall algorithm: we are currently looking at the values $X^h_{i,j}$. We will maintain four invariants:

**(I1)** Once a variable $x_{i,j}$ gets assigned to 0 or 1, it is never changed;

**(I2)** The constraints (Assign) always hold; and

**(I3)** Once a constraint in (Capacity) becomes tight, it remains tight.

**(IV)** Once a constraint is dropped in some iteration, it is never reinstated.

Iteration $(h+1)$ of **Sched-Cap** consists of three main steps:

1. Since we aim to maintain (**I1**), let us remove all $X^h_{i,j} \in \{0, 1\}$; i.e., we project $X^h$ to those co-ordinates $(i, j)$ for which $X^h_{i,j} \in (0, 1)$, to obtain the current vector $Y$ of "floating" (to-be-rounded) variables; let $\mathcal{S} \equiv (A_h Y = u_h)$ denote the current linear system that represents **LP-Cap**. ($A_h$ is some matrix and $u_h$ is a vector; we avoid using "$\mathcal{S}_h$" to simplify notation.) In particular, the "capacity" of machine $i$ in $\mathcal{S}$ is its residual capacity $b'_i$, i.e., $b_i$ minus the number of jobs that have been permanently assigned to $i$ thus far.

2. Let $Y \in \Re^v$ for some $v$; note that $Y \in (0, 1)^v$. Let $M_k$ denote the set of all machines $i$ for which exactly $k$ of the values $Y_{i,j}$ are positive. We will now drop some of the constraints in $\mathcal{S}$:

**(D1)** for each $i \in M_1$, we drop its load and capacity constraints from $\mathcal{S}$;

**(D2)** for each $i \in M_2$, we drop its load constraint and rewrite its capacity constraint as $x_{i,j_1} + x_{i,j_2} \leq \lceil X^h_{i,j_1} + X^h_{i,j_2} \rceil$, where $j_1, j_2$ are the two jobs fractionally assigned to $i$.

**(D3)** for each $i \in (M_3)$ for which *both* its load and capacity constraints are tight in $\mathcal{S}$, we drop its load constraint from $\mathcal{S}$.

3. Let $\mathcal{P}$ denote the polytope defined by this reduced system of constraints. A key claim that is proven in Lemma 3 below is that $Y$ is *not* a vertex of $\mathcal{P}$. We now invoke **RandMove**$(Y, \mathcal{P})$; this is allowable if $Y$ is indeed not a vertex of $\mathcal{P}$.

The above three steps complete iteration $(h+1)$.

It is not hard to verify that the invariants (**I1**)-(**I4**) hold true (though the fact that we drop the all-important capacity constraint for machines $i \in M_1$ may look bothersome, a moment's reflection shows that such a machine cannot have a tight capacity-constraint since its sole relevant job $j$ has value $Y_{i,j} \in (0, 1)$). Since we make at least one further constraint tight via **RandMove** in each iteration, invariant (**I4**) shows that we terminate, and that the number of iterations is at most the initial number of constraints. Let us next present Lemma 3, a key lemma:

LEMMA 3 *In no iteration is $Y$ a vertex of the current polytope $\mathcal{P}$.*

**Proof.** Suppose that in a particular iteration, $Y$ **is** a vertex of $\mathcal{P}$. Fix the notation $v$, $M_k$ etc. w.r.t. this iteration; let $m_k = |M_k|$, and let $n'$ denote the remaining number of jobs that are yet to be assigned permanently to a machine. Let us lower- and upper-bound the number of variables $v$. On the one hand, we have $v = \sum_{k \geq 1} k \cdot m_k$, by definition of the sets $M_k$; since each remaining job $j$ contributes at least two variables (co-ordinates for $Y$), we also have $v \geq 2n'$. Thus we get

$$v \geq n' + \sum_{k \geq 1} (k/2) \cdot m_k. \tag{1}$$

On the other hand, since $Y$ has been assumed to be a vertex of $\mathcal{P}$, the number $t$ of constraints in $\mathcal{P}$ that are satisfied *tightly* by $Y$, must be at least $v$. How large can $t$ be? Each current job contributes one (Assign) constraint to $t$; by our "dropping constraints" steps (**D1**), (**D2**) and (**D3**) above, the number of tight constraints ("load" and/or "capacity") contributed by the machines is at most $m_2 + m_3 + \sum_{k \geq 4} 2m_k$. Thus we have

$$v \leq t \leq n' + m_2 + m_3 + \sum_{k \geq 4} 2m_k. \tag{2}$$

Comparison of (1) and (2) and a moment's reflection shows that such a situation is possible only if: (i) $m_1 = m_3 = 0$ and $m_5 = m_6 = \cdots = 0$; (ii) the capacity constraints are tight for all machines in $M_2 \cup M_4$ – i.e., for all machines; and (iii) $t = v$.

However, in such a situation, the $t$ constraints in $\mathcal{P}$ constitute the *tight* assignment constraints for the jobs and the *tight* capacity constraints for the machines, and are hence linearly dependent (since the total assignment "emanating from" the jobs must equal the total assignment "arriving into" the machines). Thus we reach a contradiction, and hence $Y$ is not a vertex of $\mathcal{P}$. □

We next show that the final makespan is at most $2T$ with probability one:

LEMMA 4 *Let $X$ denote the final rounded vector. Algorithm **Sched-Cap** returns a schedule, where with probability one: (i) all capacity-constraints on the machines are satisfied, and (ii) for all $i$, $\sum_{j \in J} X_{i,j} p_{i,j} < \sum_j x^*_{i,j} p_{i,j} + \max_{j \in J:\, x^*_{i,j} \in (0,1)} p_{i,j}$.*

**Proof.** Part (i) mostly follows from the fact that we never drop any capacity constraint; the only care to be taken is for machines $i$ that end up in $M_1$ and hence have their capacity-constraint dropped. However, as argued soon after the description of the three steps of an iteration, note that such a machine cannot have a tight capacity-constraint when such a constraint was dropped; hence, even if the remaining job $j$ got assigned finally to $i$, its capacity constraint cannot be violated.

Let us now prove (ii). Fix a machine $i$. If at all its load-constraint was dropped, it must be when $i$ ended up in $M_1, M_2$ or $M_3$. The case of $M_1$ is argued as in the previous paragraph. So suppose $i \in M_\ell$ for some $\ell \in \{2, 3\}$ when its load constraint got dropped. Let us first consider the case $\ell = 2$. Let the two jobs fractionally assigned on $i$ at that point have processing times $(p_1, p_2)$ and fractional assignments $(y_1, y_2)$ on $i$, where $0 \leq p_1, p_2 \leq T$, and $0 < y_1, y_2 < 1$. If $y_1 + y_2 \leq 1$, we know that at the end, the assignment vector $X$ will have at most one of $X_1$ and $X_2$ being one. Simple algebra now shows that $p_1 X_1 + p_2 X_2 < p_1 y_1 + p_2 y_2 + \max\{p_1, p_2\}$ as required. If $1 < y_1 + y_2 \leq 2$, then both $X_1$ and $X_2$ can be assigned and again, $p_1 X_1 + p_2 X_2 < p_1 y_1 + p_2 y_2 + \max\{p_1, p_2\}$. For the case $\ell = 3$, we know from (**I3**) and (**D3**) that its capacity-constraint must be *tight* at some integral value $u$ at that point, and that this capacity-constraint was preserved until the end. We must have $c = 1$ or 2 here. Let us just consider the case $c = 2$; the case of $c = 1$ is similar to the case of $\ell = 2$ with $y_1 + y_2 \leq 1$. Here again, simple algebra yields that if $0 \leq p_1, p_2, p_3 \leq T$ and $0 < y_1, y_2, y_3 < 1$ with $y_1 + y_2 + y_3 = c = 2$, then for any binary vector $(X_1, X_2, X_3)$ of Hamming weight $c = 2$, $p_1 X_1 + p_2 X_2 + p_3 X_3 < p_1 y_1 + p_2 y_2 + p_3 y_3 + \max\{p_1, p_2, p_3\}$. □

Finally we have the following lemma.

LEMMA 5 *Algorithm **Sched-Cap** can be derandomized to create a schedule of cost at most $C$.*

**Proof.** Let $X^h_{i,j}$ denote the value of $x_{i,j}$ at iteration $h$. We know for all $i, j, h$, $E[X^h_{i,j}] = x^*_{i,j}$, where $x^*_{i,j}$ is solution of **LP-Cap**. Therefore, at the end, we have that the total expected cost incurred is $C$. The procedure can be derandomized directly by the method of conditional expectation, giving an 1-approximation to cost. □

Lemmas 4 and 5 yield Theorem 2.

## 3 Scheduling with Outliers

In this section, we consider GAP with outliers and with a hard profit constraint [29]. Formally, the problem is as follows, where $x_{i,j}$ is the indicator variable for job $j$ to be scheduled on machine $i$. Given $m$ machines and $n$ jobs, where job $j$ requires processing time of $p_{i,j}$ in machine $i$, incurs a cost of $c_{i,j}$ if assigned to $i$ and provides a profit of $\pi_j$ if scheduled, the goal is to minimize the makespan, $T = \max_i \sum_j x_{i,j} p_{i,j}$, subject to the constraint that the total cost $\sum_{i,j} x_{i,j} c_{i,j}$ is at most $C$ and total profit $\sum_j \pi_j \sum_i x_{i,j}$ is at least $\Pi$.

Our main contribution here is the following:

THEOREM 6 *For any constant $\epsilon > 0$, there is an efficient algorithm **Sched-Outlier** that returns a schedule of profit at least $\Pi$, cost at most $C(1+\epsilon)$ and makespan at most $(2+\epsilon)T$, where $T$ is the optimal makespan with cost $C$ and profit $\Pi$.*

Note that this is an improvement over the work of [29], that constructs a schedule with makespan $3T$ with profit $\Pi$ and cost $C(1+\epsilon)$. In addition, our approach also accommodates *fairness*, a basic requirement in dealing with outliers, especially when problems have to be run repeatedly. We formulate fairness via stochastic programs that specify

for each job $j$, a lower-bound $r_j$ on the probability that it gets scheduled. We adapt our approach to honor such requirements:

THEOREM 7 *There is an efficient randomized algorithm that returns a schedule of profit at least* $\Pi$*, expected cost at most* $2C$ *and makespan at most* $3T$ *and guarantees that for each job* $j$*, it is scheduled with probability* $r_j$*, where* $T$ *is the optimal expected makespan with expected cost* $C$ *and expected profit* $\Pi$*. If the fairness guarantee on* any one *job can be relaxed, then for every fixed* $\epsilon > 0$*, there is an efficient algorithm to construct a schedule that has profit at least* $\Pi$*, expected cost at most* $C(1 + 1/\epsilon)$ *and makespan at most* $(2 + \epsilon)T$*.*

We defer the proof of Theorem 7 to the full version, and focus on Theorem 6. Guess the optimum makespan $T$ by binary search as in [34]. If $p_{i,j} > T$, $x_{i,j}$ is set to 0. We guess all assignments $(i, j)$ where $c_{i,j} > \epsilon' C$, with $\epsilon' = \epsilon^2$. Any valid schedule can schedule at most $1/\epsilon'$ pairs with assignment costs higher than $\epsilon' C$; since $\epsilon'$ is a constant, this guessing can be done in polynomial time. For all $(i, j)$ with $c_{i,j} > \epsilon' C$, let $\mathcal{G}_{i,j} \in \{0, 1\}$ be a correct guessed assignment. The solution to the following integer program then gives an optimal solution:

$$\sum_{i,j} c_{i,j} x_{i,j} \leq C \qquad \text{(Cost)}$$

$$\sum_{i} x_{i,j} = y_j \quad \forall j \qquad \text{(Assign)}$$

$$\sum_{j} p_{i,j} x_{i,j} \leq T \quad \forall i \qquad \text{(Load)}$$

$$\sum_{j} \pi_j y_j \geq \Pi \qquad \text{(Profit)}$$

$$x_{i,j} \in \{0, 1\} \ \forall i, j; \ y_j \in \{0, 1\} \ \forall j$$

$$x_{i,j} = 0 \ \text{ if } p_{i,j} > T$$

$$x_{i,j} = \mathcal{G}_{i,j} \ \forall (i, j) \text{ such that } c_{i,j} > \epsilon' C$$

We relax the constraint "$x_{i,j} \in \{0, 1\}$ and $y_j \in \{0, 1\}$" to "$x_{i,j} \in [0, 1]$ and $y_j \in [0, 1]$" to obtain the LP relaxation **LP-Out**. We solve the LP to obtain an optimal LP solution $x^*, y^*$; we next show how **Sched-Outlier** rounds $x^*, y^*$ to obtain the claimed approximation.

Note that $x^*_{i,j} \in [0, 1]$ denotes the fraction of job $j$ assigned to machine $i$ in $x^*$. Initially, $\sum_i x^*_{i,j} = y^*_j$. Initialize $X = x^*$. The algorithm is composed of several iterations; the random values at the end of iteration $h$ of the overall algorithm are denoted by $X^h$. (Since $y_j$ is given by the equality $\sum_i x_{i,j}$, $X^h$ is effectively the set of variables.) Each iteration $h$ (except perhaps the last one) conducts a randomized update using **RandMove** on a suitable polytope constructed from a *subset* of the constraints of **LP-Out**. Therefore, for all $h$ except perhaps the last, we have $\mathsf{E}\left[X^h_{i,j}\right] = x^*_{i,j}$. A variable $X^h_{i,j}$ is said to be *floating* if it lies in $(0, 1)$, and a job is *floating* if it is not yet finally assigned. The subgraph of $(J, M, E)$ composed of the floating edges $(i, j)$, naturally suggests the following notation at any point of time: machines of "degree" $k$ in an iteration are those with exactly $k$ floating jobs assigned fractionally, and jobs of "degree" $k$ are those assigned fractionally to exactly $k$ machines in iteration $h$. Note that since we allow $y_j < 1$, there can exist singleton (i.e., degree-1) jobs which are floating.

Suppose we are at the beginning of some iteration $(h+1)$ of the overall algorithm; so we are currently looking at the values $X^h_{i,j}$. We will maintain the following invariants:

**(I1')** Once a variable $x_{i,j}$ gets assigned to 0 or 1, it is never changed;

**(I2')** If $j$ is not a singleton, then $\sum_i x_{i,j}$ remains at its initial value;

**(I3')** The constraint (Profit) always holds;

**(I4')** Once a constraint is dropped, it is never reinstated.

Algorithm **Sched-Outlier** starts by initializing with **LP-Out**. Iteration $(h + 1)$ consists of four major steps.

1. Since we aim to maintain **(I1')**, we remove all $X^h_{i,j} \in \{0, 1\}$; i.e., we project $X^h$ to those coordinates $(i, j)$ for which $X^h_{i,j} \in (0, 1)$, to obtain the current vector $Z$ of "floating" variables; let $\mathcal{S} \equiv (A_h Z = u_h)$ denote the current linear system that represents **LP-Out**. ($A_h$ is some matrix and $u_h$ is a vector.)

2. Let $Z \in \Re^v$ for some $v$; note that $Z \in (0, 1)^v$. Let $M_k$ and $N_k$ denote the set of degree-$k$ machines and degree-$k$ jobs respectively, with $m_k = |M_k|$ and $n_k = |N_k|$. We will now drop/replace some of the constraints in $\mathcal{S}$:

**(D1')** for each $i \in M_1$, we drop its load constraint from $\mathcal{S}$;

**(D2')** for each $i \in N_1$, we drop its assignment constraint from $\mathcal{S}$; we include one profit constraint, $\sum_{j \in N_1} Z_{i,j}\pi_j = \sum_{j \in N_1} X^h_{i,j}\pi_j$ that replaces the constraint (Profit). (Note that at this point, the values $X^h_{i,j}$ are some constants.)

Thus, the assignment constraints of the singleton jobs are replaced by *one* profit constraint.

3. If $Z$ is a vertex of $\mathcal{S}$ then define fractional assignment of a machine $i$ by $h_i = \sum_{j \in J} Z_{i,j}$. Define a job $j$ to be tight, if $\sum_{i \in M} Z_{i,j} = 1$. Drop all assignment constraints of the non-tight jobs and maintain a single profit constraint, $\sum_{j \in N_1 \cup J_N} Z_{i,j}\pi_j = \sum_{j \in N_1 \cup J_N} X^h_{i,j}\pi_j$, where $J_N$ are the nontight jobs. If $Z$ is not yet a vertex of the modified $\mathcal{S}$, then while there exists a machine $i'$ whose degree $d$ satisfies $h_{i'} \geq (d - 1 - \epsilon)$, drop the load constraint on machine $i'$.

4. Let $\mathcal{P}$ denote the polytope defined by this reduced system of constraints. If $Z$ is not a vertex of $\mathcal{P}$, invoke **RandMove**$(Z, \mathcal{P})$. Else, if there is a degree-3 machine with 1 singleton job, assign the singleton job and the cheaper (less processing time) of the non-singleton jobs to it. If there exist two singleton jobs, then discard the one with less profit. Assign remaining of the jobs in a way, so that each machine gets at most one extra job.

We now prove a key lemma, which shows when in step 4, $Z$ can possibly be a vertex of the polytope under consideration. Recall that in the bipartite graph $G = (J, M, E)$, we have in iteration $(h + 1)$ that $(i, j) \in E$ iff $X^h_{i,j} \in (0, 1)$. Any job or machine having degree 0 is not part of $G$.

LEMMA 8 *Let $m$ denote the number of machine-nodes in $G$. If $m \geq \frac{1}{\epsilon}$, then $Z$ is not a vertex of the current polytope.*

**Proof.** Let us consider the different possible configurations of $G$, when $Z$ becomes a vertex of the polytope $\mathcal{P}$ at step 3. There are several cases to consider depending on the number of singleton floating jobs in $G$ in that iteration.

Case 1: There is no singleton job: We have $n_1 = 0$. Then, the number of constraints in $\mathcal{S}$ is $EQ = \sum_{k \geq 2} m_k + \sum_{k \geq 2} n_k$. Also the number of floating variables is $v = \sum_{k \geq 2} kn_k$. Alternatively, $v = \sum_{k \geq 1} km_k$. Therefore, $v = \sum_{k \geq 2} \frac{k}{2}(m_k + n_k) + \frac{m_1}{2}$. $Z$ being a vertex of $\mathcal{P}$, $v \leq EQ$. Thus, we must have, $n_k, m_k = 0, \ \forall k \geq 3$ and $m_1 = 0$. Hence every floating machine has exactly two floating jobs assigned to it and every floating job is assigned exactly to two floating machines (Config-1 of Figure 1).

Case 2: There are at least 3 singleton jobs: We have $n_1 \geq 3$. Then the number of linear constraints is $EQ = \sum_{k \geq 2} m_k + \sum_{k \geq 2} n_k + 1$. The last "1" comes from considering one profit constraint for the singleton jobs. The number of variables, $v = \frac{n_1}{2} + \sum_{k \geq 2} \frac{k}{2}(m_k + n_k) + \frac{m_1}{2} \geq \frac{3}{2} + \sum_{k \geq 2} \frac{k}{2}(m_k + n_k) + \frac{m_1}{2}$. Hence the system is always underdetermined and $Z$ cannot be a vertex of $\mathcal{P}$.

Case 3: There are exactly 2 singleton jobs: We have $n_1 = 2$. Following similar counting arguments, we can show that each machine must have exactly two floating jobs assigned to it and each job except two is assigned to exactly two machines fractionally (Config-2 of Figure 1).

Case 4: There is exactly 1 singleton job: We have $n_1 = 1$. Then $EQ = \sum_{k \geq 2} m_k + \sum_{k \geq 2} n_k + 1$ and $v \geq \frac{1}{2} + n_2 + \frac{3}{2}n_3 + \frac{m_1}{2} + m_2 + \frac{3}{2}m_3 + \sum_{k \geq 4} \frac{k}{2}(m_k + n_k)$. If $Z$ is a vertex of $\mathcal{P}$, then $v \leq EQ$. There are few possible configurations that might arise in this case.

(i) Only one job of degree 3 and one job of degree 1. All the other jobs have degree 2 and all the machines have degree 2. We call this Config-3.

(ii) Only one machine of degree 3 and one job of degree 1. The rest of the jobs and machines have degree 2. We call this Config-4.

(iii) Only one machine of degree 1 and one job of degree 1. The rest of the jobs and machines have degree 2. We call this Config-5.

These configurations are shown in Fig. 1. Each configuration can have arbitrary number of disjoint cycles.

In any configuration, if there is a cycle with all tight jobs, then there always exists a machine with total fractional assignment 1 and hence its load constraint can always be dropped to make the system underdetermined. So we assume there is no such cycle in any configurations. Now suppose the algorithm reaches Config-1. If there are two non-tight jobs, then the system becomes underdetermined. Therefore, there can be at most one non-

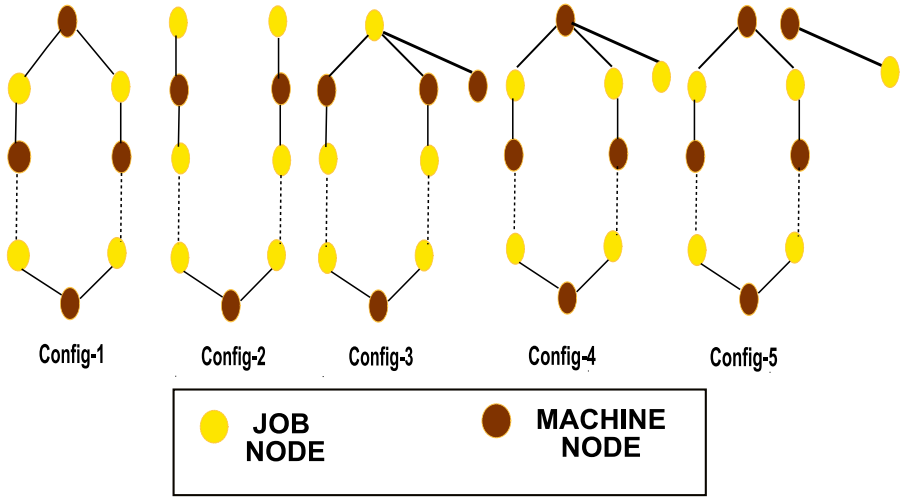


Figure 1: Different configurations of machine-job bipartite graph at Step 3 and 4

tight job and only one cycle (say $C$) with that non-tight job. Let $C$ have $m$ machines and thus $m$ jobs. Therefore, $\sum_{i,j \in C} x_{i,j} \geq m-1$. Thus there exists a machine, such that the total fractional assignment of jobs on that machine is $\geq \frac{m-1}{m} = 1 - 1/m$. If $m \geq \frac{1}{\epsilon}$, then there exists a machine with total fractional assignment $\geq (1-\epsilon)$. Dropping the load constraint on that machine makes the system underdetermined.

If the algorithm reaches Config-2, then all the jobs must be tight for $Z$ to be a vertex. If there are $m$ machines, then the number of non-singleton jobs is $m-1$. If $x_{i_1,j_1} + x_{i_2,j_2} \geq 1$, then following similar averaging argument as in Config-1, we can show the machine with maximum fractional job assignment, must have a total fractional assignment at least 1. Otherwise, if $m \geq \frac{1}{\epsilon}$, again the machine with maximum fractional job assignment, must have a total fractional assignment at least $1-\epsilon$. For Config-3 and 5, if $Z$ is a vertex of $\mathcal{P}$, then all jobs must be tight and using same argument, there exists a machine with fractional assignment at least $(1-\epsilon)$ if the algorithm reaches Config -3 and there exists a machine with fractional assignment 1, if the algorithm reaches Config-5.

If the algorithm reaches Config-4, then again all jobs must be tight. If the degree-3 machine has fractional assignment at least $2-\epsilon$, then its load constraint can be dropped to make the system underdetermined. Otherwise, the total assignment to the degree-2 machines from all the jobs in the cycle is at least $m-2+\epsilon$. Therefore, there exists at least one degree-2 machine with fractional assignment at least $\frac{m-2+\epsilon}{m-1} = 1 - \frac{1-\epsilon}{m-1} \geq 1-\epsilon$, if $m \geq \frac{1}{\epsilon}$. This completes the proof. □

We next show that the final profit is at least $\Pi$ and the final makespan is at most $(2+\epsilon)T$:

LEMMA 9 *Let $X$ denote the final rounded vector. Algorithm **Sched-Outlier** returns a schedule, where with probability one, (i) profit is at least $\Pi$, (ii) for all $i$, $\sum_{j \in J} X_{i,j} p_{i,j} < \sum_j x^*_{i,j} p_{i,j} + (1+\epsilon) max_{j \in J:\ x^*_{i,j} \in \{0,1\}} p_{i,j}$.*

**Proof.** (i) This essentially follows from the fact that whenever assignment constraint on any job is dropped, its profit constraint is included in the global profit constraint of the system. At step 4, except for one case (Config-2), all the jobs are always assigned, so profit can not decrease in those cases. A singleton job (say $j_1$) is dropped, only when $G$ has two singleton jobs $j_1, j_2$ fractionally assigned to $i_1$ and $i_2$ respectively, with total assignment $x_{i_1,j_1} + x_{i_2,j_2} < 1$. Otherwise the system remains underdetermined from Lemma 8. Since the job with higher profit is retained, $\pi_{j_1} x_{i_1,j_1} + \pi_{j_2} x_{i_2,j_2} \leq max\{\pi_{j_1}, \pi_{j_2}\}$.

(ii) From Lemma 8 and **(D1')**, load constraints are dropped from machines $i \in M_1$ and might be dropped from machine $i \in M_2 \cup M_3$. For $i \in M_1$, only the remaining job $j$ with $X^h_{i,j} > 0$, can get fully assigned to it. Hence for $i \in M_1$, its total load is bounded by $\sum_j x^*_{i,j} p_{i,j} + \max_{j \in J: x^*_{i,j} \in \{0,1\}} p_{i,j}$. For any machine $i \in M_2 \cup M_3$, if their degree $d$ (2 or 3) is such that, its fractional assignment is at least $d - 1 - \epsilon$, then by simple algebra, it can be shown that for any such machine $i$, its total load is at most $\sum_j x^*_{i,j} p_{i,j} + (1+\epsilon)\max_{j \in J: x^*_{i,j} \in \{0,1\}} p_{i,j}$. For the remaining machines consider what happens at step 4. Except when Config-4 is reached, any remaining machine $i$ gets at most one extra job, and thus its total load is bounded by $\sum_j x^*_{i,j} p_{i,j} + \max_{j \in J:\ x^*_{i,j} \in \{0,1\}} p_{i,j}$. When Config-4 is reached at step 4, if the degree-3 machine has a fractional assignment at most 1, then for any value of $m$, there will exist a degree-2 machine whose fractional assignment is 1, giving a contradiction. Hence, let $j_1, j_2, j_3$ be the three jobs assigned fractionally to the degree-3 machine $i$ and let $j_3$ be the singleton job, and $x_{i,j_1} + x_{i,j_2} > 1$. If $p_{i,j_1} \leq p_{i,j_2}$, then the degree-3 machine gets $j_1, j_3$. Else the degree-3 machine gets $j_2, j_3$. The degree-3 machine gets 2 jobs, but its fractional assignment from $j_1$ and $j_2$ is already at least 1. Since the job with less processing time among $j_1$ and $j_2$ are assigned to $i$, its increase in load can be at most $\sum_j x^*_{i,j} p_{i,j} + \max_{j \in J:\ x^*_{i,j} \in \{0,1\}} p_{i,j}$. □

Finally we have the following lemma.

LEMMA 10 *Algorithm **Sched-Outlier** can be derandomized to output a schedule of cost at most $C(1+\epsilon)$.*

**Proof.** In all iterations $h$, except the last one, for all $i, j$, $E[X^h_{i,j}] = x^*_{i,j}$, where $x^*_{i,j}$ is solution of **LP-Out**. Therefore, before the last iteration, we have that the total expected cost incurred is $C$. The procedure can be derandomized directly by the method of conditional expectation, giving an 1-approximation to cost, just before the last iteration. Now at the last iteration, since at most $\frac{1}{\epsilon}$ jobs are assigned and each assignment requires at most $\epsilon' C = \epsilon^2 C$ in cost, the total increase in cost is at most $\epsilon C$, giving the required approximation. □

Lemmas 9 and 10 yield Theorem 6.

# 4 Max-Min Fair Allocation

We now present our results for max-min fair allocation [4, 5, 8, 14, 15]. There are $m$ goods and $k$ persons. Each person $i$ has a non-negative integer valuation $u_{i,j}$ for good $j$. The valuation functions are linear, i.e. $u_{i,C} = \sum_{j \in C} u_{i,j}$ for any set of $C$ goods. The goal is to allocate each good to a person such that the "least happy person is as happy as possible": i.e., $\min_i u_{i,C}$ is maximized. Our algorithm is based upon rounding the configuration LP which is described in Subsection 4.1

THEOREM 11 *Given any feasible solution to the configuration LP, it can be rounded to a feasible integer solution such that every person gets at least $\Theta(\sqrt{\frac{k \log k}{\log \log k}})$ fraction of the optimum utility with high probability in polynomial time.*

The approximation factor of $O(\sqrt{\frac{k \log k}{\log \log k}})$ is an improvement of the previous work of [5], that achieved an approximation factor of $O(\sqrt{k} \log^3 k)$; our bound is near-optimal since the integrality gap of the configuration LP is $\Omega(\sqrt{k})$ [8]. However, note that the recent work of Chakrabarty, Chuzhoy and Khanna [20] has improved the bound to $m^\epsilon$. (Also note that $m \geq k$.) Our main point is to show the applicability of our types of rounding approaches to this variation of the problem as well.

In the context of fair allocation, an additional important criterion can be an *equitable partitioning* of goods: we may impose an upper bound on the number of items a person might receive. For example, we may want each person to receive at most $\lceil \frac{m}{k} \rceil$ goods. Theorem 1 leads to the following:

THEOREM 12 *Suppose, in max-min allocation, we are given upper bounds $c_i$ on the number of items that each person $i$ can receive, in addition to the utility values $u_{i,j}$. Let $T$ be the optimum max-min allocation value that satisfies $c_i$ for all $i$. Then, we can efficiently construct an allocation in which for each person $i$ the bound $c_i$ holds and she receives a total utility of at least $T - \max_j u_{i,j}$.*

This generalizes the result of [14], which yields the "$T - \max_j u_{i,j}$" value when no bounds such as the $c_i$ are given. To our knowledge, the results

of [4, 5, 8, 15] do not carry over to the setting of such "fairness bounds" $c_i$.

We now describe the algorithm and the proof of Theorem 11 in the next subsection. The major steps of the algorithm are similar to [5], however within each step the algorithm uses different rounding techniques; and hence our analysis is completely different. Rounding techniques are motivated by variants of the dependent rounding method.

### 4.1 Algorithm for Max-Min Fair Allocation

We start by describing *bipartite dependent rounding* [27], which is a special case of **Rand-Move** that we have discussed so far.

#### 4.1.1 Bipartite Dependent Rounding

Suppose $G = (U, V, E)$ is a bipartite graph with $U$ and $V$ being vertices in the two partitions and $E$ being the edges between them. The edges $E$ are defined by the vector $x$. If $x_{i,j} \in (0, 1), i \in U, j \in V$, then $(i, j) \in E$ and vice-versa. The dependent rounding algorithm chooses an even cycle $\mathcal{C}$ or a maximal path $\mathcal{P}$ in $G$, and partitions the edges in $\mathcal{C}$ or $\mathcal{P}$ into two matchings $\mathcal{M}_1$ and $\mathcal{M}_2$. Then, two positive scalars $\alpha$ and $\beta$ are defined as follows:

$$\begin{aligned}\alpha = \min\{\eta > 0 : ((\exists (i,j) \in \mathcal{M}_1 : x_{i,j} + \eta = 1)\\ \bigvee (\exists (i,j) \in \mathcal{M}_2 : x_{i,j} - \eta = 0))\};\\ \beta = \min\{\eta > 0 : ((\exists (i,j) \in \mathcal{M}_1 : x_{i,j} - \eta = 0)\\ \bigvee (\exists (i,j) \in \mathcal{M}_2 : x_{i,j} + \eta = 1))\};\end{aligned}$$

Now with probability $\frac{\beta}{\alpha+\beta}$, set

$$\begin{aligned} & Y_{i,j} = x_{i,j} + \alpha \text{ for all } (i,j) \in \mathcal{M}_1\\ \text{and} \quad & Y_{i,j} = x_{i,j} - \alpha \text{ for all } (i,j) \in \mathcal{M}_2;\end{aligned}$$

with complementary probability of $\frac{\alpha}{\alpha+\beta}$, set

$$\begin{aligned} & Y_{i,j} = x_{i,j} - \beta \text{ for all } (i,j) \in \mathcal{M}_1\\ \text{and} \quad & Y_{i,j} = x_{i,j} + \beta \text{ for all } (i,j) \in \mathcal{M}_2;\end{aligned}$$

The above rounding scheme satisfies the following two properties, which are easy to verify:

$$\forall \quad i,j, \quad \mathsf{E}\big[Y_{i,j}\big] = x_{i,j} \tag{3}$$

$$\exists \quad i,j, \quad Y_{i,j} \in \{0,1\} \tag{4}$$

If $X$ denotes the final rounded variable, then for any node $v \in U \bigvee V$,

$$\sum_{u \in U \bigvee V} X_{u,v} \in \{\lceil \sum_{u \in U \bigvee V} x_{u,v} \rceil, \lfloor \sum_{u \in U \bigvee V} x_{u,v} \rfloor\}. \tag{5}$$

#### 4.1.2 Algorithm

Our algorithm is based upon rounding a configuration linear program similar to [5, 8]. We guess the optimum solution- value $T$, using binary search. There is a variable $x_{i,C}$ for assigning a valid bundle $C$ to person $i$. An item $j$ is said to be small for person $i$, if $u_{i,j} < \frac{T}{\lambda}$, otherwise it is said to be big. Here $\lambda$ is the approximation ratio, which will get fixed later. A configuration is a subset of items. A configuration $C$ is called valid for person $i$, if,

- $u_{i,C} \geq T$ and all the items are small; or
- $C$ contains only one item $j$ and $u_{i,j} \geq \frac{T}{\lambda}$, that is, $j$ is a big item for person $i$.

Let $C(i, T)$ denote the set of all valid configurations corresponding to person $i$ with respect to $T$. The configuration LP relaxation of the problem is as follows:

$$\begin{aligned} \forall j : \sum_{C \ni j} \sum_{i} x_{i,C} \leq 1 \\ \forall i : \sum_{C \in C(i,T)} x_{i,C} = 1 \\ \forall i, C : x_{i,C} \geq 0 \end{aligned} \tag{6}$$

Using an argument similar to [8], one can show that if the above LP is feasible, then it is possible to find a fractional allocation that provides a bundle with value at least $(1 - \epsilon)T$ for each person in polynomial time.

We define a weighted bipartite graph $G$, with the vertex set, $A \bigcup B$ corresponding to the persons and the items respectively. There is an edge between a vertex corresponding to person $i \in A$ and item $j \in B$, if a configuration $C$ containing $j$ is fractionally assigned to $i$. Define

$$w_{i,j} = \sum_{C \ni j} x_{i,C},$$

i.e., $w_{i,j}$ is the fraction of item $j$ that is allocated to person $i$ by the fractional solution of the LP. We

know an item $j$ is big for person $i$, if $u_{i,j} \geq \frac{T}{\lambda}$. In this case, the edge $(i,j)$ is called a matching edge. Otherwise it is called a flow edge.

Let $M$ and $F$ represent the set of matching and flow edges respectively. For each vertex $v \in A \bigcup B$, let $m_v$ be the total fraction of the matching edges incident to it. Also define $f_v = 1 - m_v$. The main steps of the algorithm are shown in Table 1.

Note that these steps are similar to that in [5], but the rounding techniques within each step are different. The rounding techniques are described next.

#### 4.1.3 Finding a Random Matching

Consider the edges in $M$. Remove all the edges $(i,j)$ that have already been rounded to 0 or 1. Additionally if an edge is rounded to 1, remove both $i$ and $j$ (person $i$ is satisfied as $u_{i,j} \geq \frac{T}{\lambda}$). We know an edge $(i,j) \in M$ implies that $u_{i,j} \geq \frac{T}{\lambda}$. Therefore, even if different items have different utility, a person matched to any one item in this stage has received at least $T/\lambda$ utility. We initialize for each $(i,j) \in M$, $y_{i,j} = w_{i,j}$ and modify the $y_{i,j}$ values probabilistically in rounds using the approach described in Subsubsection 4.1.1. If $Y_{i,j}$ denotes the final rounded value, then $\forall (i,j)$ by Property (3), $\mathsf{E}\big[Y_{i,j}\big] = w_{i,j}$. This gives the following corollary.

COROLLARY 13 *The probability that a vertex $v \in A \bigcup B$ is saturated in the matching generated by the algorithm is $m_v$.*

**Proof.** Let the edges $e_1, e_2, ..e_l \in M$ are incident on $v$. Then, $\mathsf{Pr}\big[v \text{ is saturated}\big] = \mathsf{Pr}\big[\exists e_i, i \in [1,l] \;\; s.t \;\; v \text{ is matched with } e_i\big] = \sum_{i=1}^{l} \mathsf{Pr}\big[v \text{ is matched with } e_i\big] = \sum_{i=1}^{l} w_i = m_v$ □

#### 4.1.4 Allocating small bundles

Consider a person $v$, who is not saturated in the matching: how much utility does this person $v$ get ? From all the bundles which are fractionally assigned to person $v$, remove any item $j$, with $m_j \geq 1 - \epsilon_1$ ($\epsilon_1$ to be fixed later). Since the total sum of $m_j$ can be at most $k$ ($k$ = number of persons), there can be at most $\frac{k}{1-\epsilon_1}$ items in the bundles with $m_j \geq 1 - \epsilon_1$. Therefore the remaining items in the bundle have value $f_j \geq \epsilon_1$. Since bundles only contain small items, and the total valuation of each bundle (fractionally) is at least $T$, we have, after removing the items with $m_j \geq 1 - \epsilon_1$, the remaining valuation is at least $\Big(T - \frac{\epsilon_1 k}{1-\epsilon_1}\frac{T}{\lambda}\Big)$. Define a random variable $Y_{v,j}$ for each remaining item such that,

$$Y_{v,j} = \begin{cases} \frac{w'_{v,j} u_{v,j}}{T/\lambda} & \text{with probability } f_j \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

Here $w'_{v,j} = \frac{w_{v,j}}{f_v}$. Since each person $v$ is not saturated by matching with probability $1 - m_v = f_v$, each such person $v$ selects bundle $C$ with probability $x_{v,C}/f_v$. Thus each item $j$ is selected with probability $w_{v,j}/f_v = w'_{v,j}$.

Define $G_v = \sum_j Y_{v,j}$. Then $\frac{T}{\lambda} G_v$ is the total fractional assignment to each person after step (3) and after doing further processing as suggested in the beginning of this subsubsection. We have,

$$\begin{aligned} \mathsf{E}\big[G_v\big] &= \sum_j \frac{w'_{v,j} u_{v,j} f_j}{T/\lambda} \\ &\geq \epsilon_1 \lambda (1 - \frac{\epsilon_1 k}{(1-\epsilon_1)\lambda}) \end{aligned} \tag{8}$$

Now we will show in Lemma 14 that $Y_{v,j}$'s are negatively correlated. Therefore, applying Chernoff-Hoeffding bound for negatively-correlated random variables [37], we get

$$\mathsf{Pr}\big[G_v \leq (1-\epsilon)\mathsf{E}\big[G_v\big]\big] \leq exp(-E[G_v]\epsilon^2/3)$$

for any $\epsilon \in (0,1)$.

Or we have,

$$\begin{aligned} & \mathsf{Pr}\big[\frac{T}{\lambda} G_v \leq (1-\epsilon) \sum_j w'_{v,j} u_{v,j} f_j\big] \\ \leq \;\; & exp(-\mathsf{E}\big[G_v\big]\epsilon^2/3) \end{aligned}$$

Now substitute, $\epsilon_1 = \sqrt{\frac{\log k}{\log\log k}}\frac{1}{\sqrt{k}}$ and $\lambda = 2\sqrt{k}\sqrt{\frac{\log k}{\log\log k}}$. We have,

$$\begin{aligned} & \mathsf{Pr}\big[\frac{T}{\lambda} G_v \leq (1-\epsilon) \sum_j w'_{v,j} u_{v,j} f_j\big] \\ \leq \;\; & exp(-\mathsf{E}\big[G_v\big]\epsilon^2/3) \\ \leq \;\; & exp(-\frac{\log k}{2\log\log k}\epsilon^2/3) \\ = \;\; & \Theta(\frac{\log k}{k}) \end{aligned}$$

| **Main Steps of the Algorithm for Max-Min Fair Allocation** |
|---|
| (1) Guess the value of the optimum solution $T$ by doing a binary search. Solve LP (6). Obtain the set $M$ and $m_v$, $f_v$ for each vertex $v$ in $G$. |
| (2) Select a random matching from edges in $M$ using the algorithm described in Subsubsection 4.1.3, such that for every $v \in A \bigcup B$ the probability that $v$ is saturated by the matching is $m_v = 1 - f_v$. |
| (3) Let every person $i$ who is not matched yet selects a bundle $C$ of small goods with probability $\frac{x_{i,C}}{f_i}$ and claims the goods in that bundle, except those already assigned in the previous step. |
| (4) For each good $j$ claimed by several persons in the previous step, resolve the contention by following the algorithm described in Subsubsection 4.1.5. |

Table 1: High level description of the algorithm for max-min fair allocation

Therefore in this step, the net fractional utility assigned to each person is $\geq \sqrt{\frac{\log k}{k \log\log k}} \frac{T(1-\epsilon)}{2}$, with probability $\geq 1 - \Theta(\log k/k)$.

Now we prove that $Y_{v,j}$'s are negatively correlated.

LEMMA 14 *The random variables $Y_{v,j}$, $j = 1, 2, .., n$ as defined in Equation (7) are negatively correlated.*

**Proof.** Define an indicator random variable

$$X_j = \begin{cases} 1 & \text{if } j \text{ is saturated in the matching step} \\ 0 & \text{otherwise} \end{cases}$$

We will show that $\forall b \in \{0, 1\}$, for any subset of jobs $S$, $\mathsf{Pr}\big[\bigwedge_{j \in S}(X_j = b)\big] \leq \prod_{j \in S} \mathsf{Pr}\big[X_j = b\big]$. This will imply that the $Y_{v,j}$'s are negatively correlated.

Fix a subset of items $J$. Let $b = 1$ (the proof for $b = 0$ is identical). Consider iteration $h$. Let $H_j^h = \sum_i y_{i,j}^h$, where $y_{i,j}^h$ denote the value of $y_{i,j}$ at the beginning of iteration $h$. We will show,

$$\forall i \;\; \mathsf{E}\big[\prod_{j \in J} H_j^h\big] \leq \mathsf{E}\big[\prod_{j \in J} H_j^{(h-1)}\big] \qquad (9)$$

Thus we will have,

$$\begin{aligned} & \mathsf{Pr}\big[\bigwedge_{j \in J}(X_j = 1)\big] \\ = \;& \mathsf{E}\big[\prod_{j \in J} H_j^{|E|+1}\big] \leq \mathsf{E}\big[\prod_{j \in J} H_j^1\big] \\ = \;& \prod_{j \in J} \sum_l y_{i,j} = \prod_{j \in J} \mathsf{Pr}\big[X_j = 1\big] \end{aligned}$$

Let us now prove (9) for a fixed $i$. In iteration $i$, exactly one of the following three cases occur:

**Case 1**: *All the jobs $j \in J$ are internal nodes of the maximal path. (If it is a cycle all the nodes are internal).*

In this case, the value of $H_j^h$'s, $j \in J$, do not change. Hence, $\mathsf{E}\big[\prod_{j \in J} H_j^h | \text{Case 1}\big] \leq \mathsf{E}\big[\prod_{j \in J} H_j^{(h-1)}\big]$.

**Case 2**: *Exactly one job, say $j_1 \in J$ is the end point of the maximal path considered in iteration $h$, or has its value modified.*

Let $B(j_1, \alpha, \beta)$ denote the event that the job $j_1$ has its value modified in the following probabilistic way:

$$H_{j_1}^h = \begin{cases} H_{j_1}^{h-1} + \alpha & \text{with probability } \frac{\beta}{\alpha+\beta} \\ H_{j_1}^{h-1} - \beta & \text{with probability } \frac{\alpha}{\alpha+\beta} \end{cases}$$

Thus,

$$\mathsf{E}\big[H_{j_1}^h | \forall j \in J,\; H_j^{h-1} = a_j \wedge B(j_1, \alpha, \beta)\big] = a_{j_1}$$

Since the values of $H_j$, $j \neq j_1$ remains unchanged and the above equation holds for any $j_1, \alpha, \beta$, we have the desired result.

**Case 3**: *Two jobs, say $j_1$ and $j_2$ are the end points of the maximal path considered in iteration $i$, or have their values modified.*

See Event A of Lemma 2.2 of [27]. □

### 4.1.5 Contention Resolution

Consider the subgraph of the flow-graph in which an edge between a person and an item

remains if and only if the item is claimed by the person in the previous step. We showed each person has a net fractional utility of at least $\frac{T}{2}\sqrt{\frac{\log k}{k \log\log k}}(1-\epsilon) = \Theta(\frac{T\sqrt{\log k}}{\sqrt{k}\log\log k})$ in this subgraph. The weight on an edge between person $v$ and item $j$ in this subgraph is $w'_{v,j}$ and the utility of an item $j$ to person $v$ is $u_{v,j}$. Now we again do a kind of dependent rounding on this subgraph, where we additionally consider the utility of the items while modifying the assignment values on the edges. This is partly motivated by [46]. We remove all $(i,j)$ that have already been rounded to 0 or 1. Let $F'$ be the current graph consisting of those $w'_{i,j}$ that lie in $(0,1)$. Choose any maximal path $P = (v_0, v_1, .., v_s)$ or a cycle $C = (v_0, v_1, .., v_s = v_0)$. The current $w'$ value of an edge $e_t = (v_{t-1}, v_t)$ is denoted by $y_t$, that is $y_t = w'_{t-1,t}$.

We will next choose the values $z_1, z_2, .., z_s$ either deterministically or probabilistically, depending on whether a cycle or a maximal path is chosen. We will update the $w'$ value of each edge $e_t = (v_{t-1}, v_t)$ to $y_t + z_t$.

Suppose we have initialized some value for $z_1$ and that we have chosen the increments $z_1, z_2, \ldots, z_t$ for some $t \geq 1$. Then the value $z_{t+1}$ corresponding to the edge $e_{t+1} = (v_t, v_{t+1})$ is chosen as follows:

**(PI)** $v_t$ is an item, then $v_{t+1} = -v_t$. (Each item is not assigned more than once.)

**(PII)** $v_t$ is a person. Then choose $z_{t+1}$ so that the utility of $w_t$ remains unchanged. Set $z_{t+1} = \frac{-u_{v_t,v_{t-1}}}{u_{v_t,v_{t+1}}}$

The vector $z = (z_1, z_2, ..z_s)$ is completely determined by $z_1$. We denote this by $f(z)$.

Now let $\mu$ be the smallest positive value such that if we get $z_1 = \mu$, then all the $w'$ values (after incrementing by the vector $z$ as specified above) stay in $[0,1]$, and at least one of them becomes 0 or 1. Similarly, let $\gamma$ be the smallest value such that if we set $z_1 = -\gamma$, then this rounding progress property holds. Now when we are considering, a maximal path, we choose the vector $z$ as follows:

**(RPI)** with probability $\frac{\gamma}{\mu+\gamma}$, let $z = f(\mu)$;

**(RPII)** with the complementary probability of $\frac{\mu}{\mu+\gamma}$, let $z = f(-\gamma)$.

Therefore in this case, if $Z = (Z_1, Z_2, ..Z_s)$ denote the random vector $z$ chosen in steps **(RPI)** and **(RPII)**, the choice of probabilities in **(RPI)** and **(RPII)** ensures that $E[Z_1] = 0$, and since the rest are functions of $z_1$ alone, $\mathsf{E}\big[Z_t\big] = 0$ for all $t$.

Now when we are considering a cycle, assume $v_0$ is a person. The assignment of $z_i$ values ensure all the objects in the cycle are assigned exactly once and utility of all the persons except $v_0$ remains unaffected. Now the change in the value of $z_s$ is $-z_1 \frac{u_{v_2,v_1} u_{v_4,v_3} \ldots u_{v_{s-1},v_{s-2}}}{u_{v_2,v_3} u_{v_4,v_5} \ldots u_{v_{s-1},v_s}}$. If $\frac{u_{v_2,v_1} u_{v_4,v_3} \ldots u_{v_{s-1},v_{s-2}}}{u_{v_2,v_3} u_{v_4,v_5} \ldots u_{v_{s-1},v_s}} > 1$, we set $z_1 = -\gamma$, else we set $z_1 = \mu$. Therefore the utility of the person $v_0$ can only increase.

Let $Y_v^i$ denote the utility assigned to person $v$ (fractional and integral) at the end of iteration $i$. The value $Y_v^0$ refers to the initial utilities in the flow-graph. Property **(PII)** and deterministic rounding scheme while considering a cycle ensures that as long as a person has degree 2 in the flow-graph $Y_v^i \geq Y_v^0$ with probability 1. In particular if $v$ never has degree 1, then its final utility is same as its initial utility in the flow graph. Suppose the degree of person $v$ becomes 1 at some iteration $i$ and let $j$ be its unique neighbor. Let $\beta = u_{v,j}$ and suppose, at the end of the iteration $i$, the total already rounded utility on person $v$ and the value of $w'_{v,j}$ are $\alpha \geq 0$ and $p \in (0,1)$ respectively. Note that $j, \alpha, \beta, p$ are all random variables and that $Y_v^i = \alpha + \beta p$; so,

$$\mathsf{Pr}\big[\alpha + \beta p \geq Y_v^0\big] = 1$$

Fix any $j, \alpha, \beta, p$ such that $\alpha + \beta p \geq Y_v^0$. Induction on the iterations show that the final utility of $v$ is $\alpha$ with probability $(1-p)$ and $\alpha + \beta$ with probability $p$. Thus the expected utility is $\alpha + \beta p$, which is same as the initial utility of $\frac{T}{2}\sqrt{\frac{\log k}{k \log\log k}}$.

In this process, there are some deterministic rounding steps interleaved with randomized rounding. We can ignore the deterministic rounding steps, since they always increase the utility. Define $X_{v,j} = 1$ if $w'_{v,j} > 0$ and $j$ was given to $v$, else define it to 0. We can prove similar to Lemma 14 that the variables $X_{v,j}$'s are negatively correlated. Now since utility of each item is at most $\frac{T}{\lambda}$, using the Chernoff-Hoeffding bounds for negative correlation [37], we get that the net utility is concentrated around its expected value with probability $\geq 1 - exp(-\frac{\frac{T}{2}\sqrt{\frac{\log k}{k \log\log k}}}{T/\lambda}) > 1 - \frac{\log k}{k}$.

Therefore we get Theorem 11.

## 5 Designing Overlay Multicast Networks For Streaming

The work of [2] studies approximation algorithms for designing a multicast overlay network. We first describe the problem and state the results in [2] (Lemma 15 and Lemma 16). Next, we show our main improvement in Lemma 17.

The background text here is largely borrowed from [2]. An overlay network can be represented as a tripartite digraph $N = (V, E)$. The nodes $V$ are partitioned into sets of entry points called sources ($S$), reflectors ($R$), and edge-servers or sinks ($D$). There are multiple commodities or streams, that must be routed from sources, via reflectors, to the sinks that are designated to serve that stream to end-users. Without loss of generality, we can assume that each source holds a single stream. Now given a set of streams and their respective edge-server destinations, a cheapest possible overlay network must be constructed subject to certain *capacity*, *quality*, and *reliability* requirements. There is a cost associated with usage of every link and reflector. There are capacity constraints, especially on the reflectors, that dictate the maximum total bandwidth (in bits/sec) that the reflector is allowed to send. The quality of a stream is directely related to whether or not an edge-server is able to reconstruct the stream without significant loss of accuracy. Therefore even though there is some loss threshold associated with each stream, at each edge-server only a maximum possible reconstruction-loss is allowed. To ensure reliability, multiple copies of each stream may be sent to the designated edge-servers.

All these requirements can be captured by an integer program. Let us use indicator variable $z_i$ for building reflector $i$, $y_{i,k}$ for delivery of $k$-th stream to the $i$-th reflector and $x_{i,j,k}$ for delivering $k$-th stream to the $j$-th sink through the $i$-th reflector. $F_i$ denotes the fanout constraint for each reflector $i \in R$. Let $p_{x,y}$ denote the failure probability on any edge (source-reflector or reflector-sink). We transform the probabilities into weights: $w_{i,j,k} = -\log(p_{k,i} + p_{i,j} - p_{k,i}p_{i,j})$. Therefore, $w_{i,j,k}$ is the negative log of the probability of a commodity $k$ failing to reach sink $j$ via reflector $i$. On the other hand, if $\phi_{j,k}$ is the minimum required success probability for commodity $k$ to reach sink $j$, we instead use $W_{j,k} = -\log(1 - \phi_{j,k})$. Thus $W_{j,k}$ denotes the negative log of maximum allowed failure. $r_i$ is the cost for opening the reflector $i$ and $c_{x,y,k}$ is the cost for using the link $(x, y)$ to send commodity $k$. Thus we have the IP (see Table 2).

Constraints (10) and (11) are natural consistency requirements; constraint (12) encodes the fanout restriction. Constraint (13), the *weight* constraint, ensures quality and reliability. Constraint (14) is the standard integrality-constraint that will be relaxed to construct the LP relaxation.

There is an important stability requirement that is referred as *color constraint* in [2]. Reflectors are grouped into $m$ color classes, $R = R_1 \cup R_2 \cup \ldots \cup R_m$. We want each group of reflectors to deliver not more than one copy of a stream into a sink. This constraint translates to

$$\sum_{i \in R_l} x_{i,j,k} \leq 1 \; \forall j \in D, \; \forall k \in S, \; \forall l \in [m] \quad (15)$$

Each group of reflectors can be thought to belong to the same ISP. Thus we want to make sure that a client is served only with one – the best – stream possible from a certain ISP. This diversifies the stream distribution over different ISPs and provides stability. If an ISP goes down, still most of the sinks will be served. We refer the LP-relaxation of integer program (Table 2) with the color constraint (15) as **LP-Color**.

All of the above is from [2].

The work of [2] uses a two-step rounding procedure and obtains the following guarantee.

First stage rounding: Rounds $z_i$ and $y_{i,k}$ for all $i$ and $k$ to decide which reflector should be open and which streams should be sent to a reflector. The results from rounding stage 1 can be summarized in the following lemma:

LEMMA 15 ***([2])*** *The first-stage rounding algorithm incurs a cost at most a factor of $64 \log |D|$ higher than the optimum cost, and with high probability violates the weight constraints by at most a factor of $\frac{1}{4}$ and the fanout constraints by at most a factor of 2. Color constraints are all satisfied.*

By incurring a factor of $\Theta(\log n)$ in the cost, the constant factors losses in the weights and fanouts can be improved.

Second stage rounding: Rounds $x_{i,j,k}$'s using the open reflectors and streams that are sent to dif-

$$\begin{aligned}
\min \quad & \sum_{i\in R} r_i z_i + \sum_{i\in R}\sum_{k\in S} c_{k,i,k} y_{i,k} + \sum_{i\in R}\sum_{k\in S}\sum_{j\in D} c_{i,j,k} x_{i,j,k} \\
s.t \quad & y_{k,i} \le z_i \ \ \forall i \in R, \ \ \forall k \in S && (10)\\
& x_{i,j,k} \le y_{i,k} \ \ \forall i \in R, \ \ \forall j \in D, \ \ \forall k \in S && (11)\\
& \sum_{k\in S}\sum_{j\in D} x_{i,j,k} \le F_i z_i \ \ \forall i \in R && (12)\\
& \sum_{i\in R} x_{i,j,k} w_{i,j,k} \ge W_{j,k} \ \ \forall j \in D, \forall k \in S && (13)\\
& x_{i,j,k} \in \{0,1\}, y_{i,k} \in \{0,1\}, z_i \in \{0,1\} && (14)
\end{aligned}$$

Table 2: Integer Program for Overlay Multicast Network Design

ferent reflectors in the first stage. The results in this stage can be summarized as follows:

LEMMA 16 ***([2])*** *The second-stage rounding incurs a cost at most a factor of* $14$ *higher than the optimum cost and violates each of fanout, color and weight constraint by at most a factor of* $7$*.*

Our main contribution is an improvement of the second-stage rounding through the use of repeated **RandMove** and by judicious choices of constraints to drop. Let us call the linear program that remains just at the end of first stage **LP-Color2**. More precisely, we show:

LEMMA 17 ***LP-Color2*** *can be efficiently rounded such that cost and weight constraints are satisfied exactly, fanout constraints are violated at most by additive* $1$ *and color constraints are violated at most by additive* $3$*.*

We defer the proof of the above lemma to the full version.

## 6 Future Directions

We discuss two speculative directions related to our rounding approach that appear promising.

Recall the notions of discrepancy and linear discrepancy from the introduction. A well-known result here, due to [12], is that if $A$ is "$t$-bounded" (every column has at most $t$ nonzeroes), then lindisc$(A) \le t$; see [31] for a closely-related result. These results have also helped in the development of improved rounding-based approximation algorithms [9, 47]. A major open question from [12] is whether lindisc$(A) \le O(\sqrt{t})$ for any $t$-bounded matrix $A$; this, if true, would be best-possible. Ingenious melding of randomized rounding, entropy-based arguments and the pigeonhole principle have helped show that lindisc$(A) \le O(\sqrt{t}\log n)$ [11, 35, 44], improved further to $O(\sqrt{t\log n})$ in [7]. However, the number of columns $n$ may not be bounded as a function of $t$, and it would be very interesting to even get some $o(t)$ bound on lindisc$(A)$, to start with. We have preliminary ideas about using the random-walks approach where the subspace $S$ (that is orthogonal to the set of tight constraints $\mathcal{C}$ in our random-walks approach) has "large" – $\Theta(n)$ – dimension. In a little more detail, whereas the constraints for rows $i$ of $A$ are dropped in [12] when there are at most $t$ to-be-rounded variables corresponding to the nonzero entries of row $i$, we propose to do this dropping at some function such as $c_0 t$ to-be-rounded variables, for a large-enough constant $c_0$ (instead of at $t$). This approach seems promising as a first step, at least for various models of random $t$-bounded matrices.

Second, there appears to be a deeper connection between various forms of dependent randomized rounding – such as ours – and iterated rounding [25, 30, 33, 42, 49]. In particular: (i) the result that we improve upon in § 2 is based on iterated rounding [18]; (ii) certain "budgeted" assignment problems that arise in keyword auctions give the same results under iterated rounding [16] and *weighted* dependent rounding [46]; and (iii) our ongoing work suggests that our random-walk ap-

proach improves upon the iterated-rounding-based work of [28] on bipartite matchings that are simultaneously "good" w.r.t. multiple linear objectives (this is related to, but not implied by, Theorem 1). We believe it would be very fruitful to understand possible deeper links between these two rounding approaches, and to develop common generalizations thereof using such insight.

## References


[1] A. Ageev and M. Sviridenko. Pipage rounding: a new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004.

[2] K. Andreev, B. Maggs, A. Meyerson, and R. Sitaraman. Designing overlay multicast networks for streaming. In *SPAA*, pages 149–158, 2003.

[3] S. Arora, A. Frieze, and H. Kaplan. A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. *Mathematical Programming*, pages 1–36, 2002.

[4] A. Asadpour, U. Feige, and A. Saberi. Santa Claus meets Hypergraph Matchings. In *Proc. APPROX-RANDOM*, pages 10–20, 2009.

[5] A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *STOC '07: Proceedings of the thirty-ninth annual ACM Symposium on Theory of computing*, pages 114–121, 2007.

[6] Y. Azar and A. Epstein. Convex programming for scheduling unrelated parallel machines. In *Proc. of the ACM Symposium on Theory of Computing*, pages 331–337. ACM, 2005.

[7] W. Banaszczyk. Balancing vectors and gaussian measures of $n$-dimensional convex bodies. *Random Structures & Algorithms*, 12:351–360, 1998.

[8] N. Bansal and M. Sviridenko. The Santa Claus problem. In *STOC '06: Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, pages 31–40, 2006.

[9] A. Bar-Noy, S. Guha, J. (Seffi) Naor, and B. Schieber. Multicasting in heterogeneous networks. In *Proc. ACM Symposium on Theory of Computing*, pages 448–453, 1998.

[10] M. Bateni, M. Charikar, and V. Guruswami. Maxmin allocation via degree lower-bounded arborescences. In *STOC '09: Proceedings of the 41st annual ACM Symposium on Theory of computing*, pages 543–552, 2009.

[11] J. Beck. Roth's estimate on the discrepancy of integer sequences is nearly sharp. *Combinatorica*, 1:319–325, 1981.

[12] J. Beck and T. Fiala. "Integer-making" theorems. *Discrete Applied Mathematics*, 3:1–8, 1981.

[13] J. Beck and V. T. Sós. *Discrepancy theory*, volume II, chapter 26, pages 1405–1446. Elsevier Science B.V. and the MIT Press, 1995.

[14] Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *SIGecom Exch.*, 5(3):11–18, 2005.

[15] D. Chakrabarty, J. Chuzhoy, and S. Khanna. On allocating goods to maximize fairness. In *FOCS '09: 50th Annual IEEE Symposium on Foundations of Computer Science*, 2009.

[16] D. Chakrabarty and G. Goel. On the Approximability of Budgeted Allocations and Improved Lower Bounds for Submodular Welfare Maximization and GAP. In *FOCS*, pages 687–696, 2008.

[17] Chandra Chekuri and Sanjeev Khanna. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM J. Comput.*, 35(3):713–728, 2005.

[18] Z. Chi, G. Wang, X. Liu, and J. Liu. Approximating scheduling machines with capacity constraints. In *FAW '09: Proceedings of the Third International Frontiers of Algorithmics Workshop*, pages 283–292, 2009. Corrected version available as arXiv:0906.3056.

[19] J. Chuzhoy and J. (Seffi) Naor. Covering problems with hard constraints. *SIAM Journal on Computing*, 36:498–515, 2006.

[20] Julia Chuzhoy and Paolo Codenotti. Resource minimization job scheduling. In *APPROX*, 2009.

[21] T. Ebenlendr, M. Křćal, and J. Sgall. Graph balancing: a special case of scheduling unrelated parallel machines. In *SODA '08: Proceedings of the Nineteenth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 483–490, 2008.

[22] M. M. Etschmaier and D. F. X. Mathaisel. Airline scheduling: An overview. *Transportation Science*, 19(2):127–138, 1985.

[23] S. Eubank, V. S. A. Kumar, M. V. Marathe, A. Srinivasan, and N. Wang. Structural and algorithmic aspects of massive social networks. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 711–720, 2004.

[24] U. Feige. On allocations that maximize fairness. In *SODA '08: Proceedings of the Nineteenth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 287–293, 2008.

[25] Lisa Fleischer, Kamal Jain, and David P. Williamson. An iterative rounding 2-approximation algorithm for the element connectivity problem. In *FOCS' 01: Proceedings of the*

*42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 339–347, 2001.

[26] R. Gandhi, E. Halperin, S. Khuller, G. Kortsarz, and A. Srinivasan. An improved approximation algorithm for vertex cover with hard capacities. *Journal of Computer and System Sciences*, 72:16–33, 2006.

[27] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM*, 53:324–360, 2006.

[28] F. Grandoni, R. Ravi, and M. Singh. Iterative rounding for multi-objective optimization problems. In *ESA '09: Proceedings of the 17th Annual European Symposium on Algorithms*, 2009.

[29] A. Gupta, R. Krishnaswamy, A. Kumar, and D. Segev. Scheduling with outliers. In *Proc. APPROX*, 2009. Full version available as arXiv:0906.2020.

[30] Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21:39–60, 2001.

[31] R. M. Karp, F. T. Leighton, R. L. Rivest, C. D. Thompson, U. V. Vazirani, and V. V. Vazirani. Global wire routing in two-dimensional arrays. *Algorithmica*, 2:113–129, 1987.

[32] V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. A unified approach to scheduling on unrelated parallel machines. *Journal of the ACM*, 56(5), 2009.

[33] Lap Chi Lau, Joseph Naor, Mohammad R. Salavatipour, and Mohit Singh. Survivable network design with degree or order constraints. In *STOC '07: Proceedings of the thirty-ninth annual ACM Symposium on Theory of computing*, pages 651–660, 2007.

[34] J. K. Lenstra, D. B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990.

[35] J. Matoušek and J. H. Spencer. Discrepancy in arithmetic progressions. *Journal of the American Mathematical Society*, 9:195–204, 1996.

[36] M. Pál, É. Tardos, and T. Wexler. Facility location with nonuniform hard capacities. In *Proc. Forty-Second Annual Symposium on Foundations of Computer Science*, pages 329–338, 2001.

[37] A. Panconesi and A. Srinivasan. Randomized distributed edge coloring via an extension of the chernoff-hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368, 1997.

[38] C. H. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 86–92, 2000.

[39] P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.

[40] R. Rushmeier, K. Hoffman, and M. Padberg. Recent advances in exact optimization of airline scheduling problems. Technical report, George Mason University, 1995.

[41] D. B. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming*, 62:461–474, 1993.

[42] Mohit Singh and Lap C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *STOC '07: Proceedings of the thirty-ninth annual ACM Symposium on Theory of computing*, pages 661–670, 2007.

[43] M. Skutella. Convex quadratic and semidefinite relaxations in scheduling. *Journal of the ACM*, 46(2):206–242, 2001.

[44] J. H. Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289:679–706, 1985.

[45] A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *IEEE Symposium on Foundations of Computer Science*, pages 588–597, 2001.

[46] A. Srinivasan. Budgeted allocations in the full-information setting. In *Proc. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 247–253, 2008.

[47] A. Srinivasan and C.-P. Teo. A constant-factor approximation algorithm for packet routing and balancing local vs. global criteria. *SIAM Journal on Computing*, 30:2051–2068, 2001.

[48] L. Tsai. Asymptotic analysis of an algorithm for balanced parallel processor scheduling. *SIAM J. Comput.*, 21(1):59–64, 1992.

[49] Vijay V. Vazirani. *Approximation Algorithms*. Springer-Verlag, 2001.

[50] G. Woeginger. A comment on scheduling two parallel machines with capacity constraints. *Discrete Optimization*, 2(3):269–272, 2005.

[51] H. Yang, Y. Ye, and J. Zhang. An approximation algorithm for scheduling two parallel machines with capacity constraints. *Discrete Appl. Math.*, 130(3):449–467, 2003.

[52] J. Zhang and Y. Ye. On the Budgeted MAX-CUT problem and its Application to the Capacitated Two-Parallel Machine Scheduling, 2001.